

APPENDIX G: WRITTEN SUBMISSIONS FROM KEY DISCUSSANTS

Expert panelists and the Workshop Co-Moderator, Dr. Roy Albert, provided the following written comments after the workshop on the questions discussed on June 23, 2000:

G.1. Comments of Dr. Roy Alpert, Division of Environmental Health, University of Cincinnati

a) A Way to Characterize Overall Carcinogen Risk: If the EPA's Air Office is going to estimate the monetary benefits of regulating carcinogenic HAPs, how will it deal with agents whose probability of being a human carcinogen is less than certain? One can estimate the number of cancer deaths by use of a dose response function and an exposure estimate on the assumption that the agent in question is a human carcinogen. But suppose the evidence does not permit us to say that it is definitely a human carcinogen? What then? One possibility is to multiply the estimated number of cancer cases by a weighting factor that is determined by the strength of the evidence. Suggested values for weighting factors are given in the table below. The use of the square of the weighting factors gives greater separation between strong and weak evidence.

EPA Category	Descriptor	Weighting Factor	(Weighting Factor) ²	Level of Evidence
A	Definite	1.0	1	Sufficient human and animal
B1	Probable	0.75	0.56	2 species positive and some human
B2	Probable	0.5	0.25	2 species positive
C	Possible	0.25	0.06	1 species both sexes
C	Possible	0.12	0.01	1 species 1 sex

b) A Method for Best-Estimate and Uncertainty Characterization of Carcinogen Risks. The stated purpose of the meeting was the development of a best-estimate and uncertainty characterization for hazard and dose-response functions for use in benefit analysis of HAP regulations. One possibility is offered below. The proposed method avoids large downward extrapolations to vanishingly small levels of risk.

The EPA's approach to carcinogen risk estimation was developed by its Carcinogen Assessment Group (CAG) in the mid 1970s. The first method involved the use of the lowest statistically significant data point as the point of departure for a linear non-threshold extrapolation to zero dose/zero incidence. Shortly after that, because of the emphasis on the conservative approach to risk assessment, the linear non-threshold extrapolation used the 95% upper confidence limit of the lowest statistically significant data point as the point of departure for the extrapolation to zero dose/zero incidence. The risk estimates obtained this way were described as "plausible upper limit estimates, i.e., ones that were not likely to be higher than the true risks but could be lower even to a considerable extent." After a few years, the statisticians objected to the procedure, saying that all the data points above the lowest statistically significant point were being wasted. After much discussion, the use of the multistage model was then agreed upon and introduced by CAG and remains the method used up to the present.

The multistage model appealed to the statisticians because it is flexible enough to accommodate almost any data set and it has biological plausibility, since cancer is multistage in its development. Furthermore, the multistage model has a low-dose linear non-threshold component. This low-dose linear non-threshold component depends on the assumption that the carcinogen in question behaves like whatever it is that causes tumors in control animals (the background). The 95% upper confidence limit for the low-dose linear non-threshold extrapolation was carried over to the multistage model. In all the variants of the linear extrapolation, the background tumor incidence was subtracted from the incidence obtained with each of the carcinogen doses so that the extrapolation could be extended to the zero dose/zero incidence. No consideration was taken of the statistical uncertainty in the background incidence.

A more realistic estimate of risk with confidence limits can be made while retaining the essential features of the established EPA approach. The multistage model can be retained on the basis of the original rationale for using it. The low-dose linear non-threshold component can also be retained on the grounds that carcinogens probably act the same way as the causes of background tumors. However, in the proposed approach we do not subtract out the background tumor incidence in order to extrapolate down to zero dose/zero incidence. We extrapolate down to the background tumor incidence. This follows logically from the assumption that the carcinogen in question acts like whatever it is that causes background tumors. The dose response curve is readjusted so that the administered dose is an increment to the equivalent background dose as determined by the slope of the linear component of the multistage model. The risk can be expressed as either the absolute incremental risk or the relative incremental risk. For example a 1% absolute risk increment on a 5% background risk would be a relative incremental risk of 20%.

There are two statistical uncertainties. The first deals with how well the data fit the extrapolation model. The second uncertainty relates to the background cancer level. The overall uncertainty of the risk estimates would be the combination of the two uncertainties.

The above approach was published in the context of time to tumor data rather than life-time incidence. But the principle is the same. The reference is as follows:

Albert, RE. and B Altschuler, 1976. Assessment of Environmental Carcinogen Risks in Terms of Lifespan Shortening. *Environmental Health Perspectives* 13: 91-94.

G.2. Comments of Dr. John C. Bailar III, Department of Health Studies, University of Chicago, Chicago, IL.

How best to identify limitations and uncertainties in both risk assessment methods and economic models.

Here I will summarize some of the important limitations and uncertainties, put them in a context, and offer some suggestions about how to deal with them.

Cost-benefit analysis is an information-hungry process, which we must apply to an information-sparse problem. This can be done, and it will be done, but the results will not be pretty.

Further, regulatory decision-making based on cost-benefit analysis is a precision-hungry process that we must necessarily base on precision-sparse inputs. Again, this can and will be done, but the results will not be pretty.

The two fundamental problems are the enormous burden of work required to deal with 189 HAPs and the great uncertainty inherent in the estimates on both sides of this process - the costs and the benefits -- for each one of them. It may be close to impossible for conscientious economists and conscientious risk assessors to do an even marginally competent job for each HAP. It may be even more difficult to remain honest about the real level of uncertainty. The more we have learned about any particular hazard, the more complex we have found it to be. This seems to be a general phenomenon, and we should adapt to it. There is no reason to think that any of the 189 HAPs is basically any simpler than benzene, though benzene looks quite complicated because we know a great deal about it.

We never think we know as much as the users of our analyses demand, and we never really know as much as we think we do. Uncertainty to three orders of magnitude is the norm in risk assessment. When that is compounded with the deep and numerous problems of cost-benefit analysis the uncertainties may very well rise to six orders of magnitude. This is in part because we need to estimate marginal effects, both costs and benefits, and these marginal effects are small differences in sometimes quite large basic numbers.

Some of the sources of uncertainty are:

- a) Missing, incomplete, and inaccurate records of human exposure; for ambient HAPs, these gaps are enormous, and we know even less about probable levels of future exposures
- b) A nearly complete absence of information about co-exposures

- c) Poor and incomplete records of human health outcomes, including the whole range of biases recognized by epidemiologists (non-random dropouts, healthy worker effects, recall biases, changing concepts and tools for diagnosis of disease, and all the rest)
- d) Unknown mechanisms of action, even when toxic endpoints have been identified, and the consequent difficulties of predicting what will happen at biologically low exposures
- e) Targets that are moving over time, including the effects on future costs and benefits of increasing life span and increasing time in which delayed toxicities can appear, decreases in competing causes of illness and death, and the different and changing medical implications of various illnesses and impairments
- f) The over-simplifications imposed by the regulatory context: Human disease rarely has one cause, and one cause does not always produce a single disease, so that we try to satisfy unifactorial regulatory demands in a highly multifactorial situation.
- g) Numerous extrapolations, which (for animal studies) include: animal to human, high to low dose, one route of administration to another, constant lifetime to intermittent and/or irregular exposure, and uniform and protected laboratory environments to highly diverse and unprotected human situations
- h) Large standard errors, sometimes from necessarily small samples, in many of the critical inputs
- i) Unpredictable and poorly understood environmental transport, including meteorology, hydrology, and many other things
- j) Possibly important synergies in biologic effects and a gross lack of understanding about two critical issues: what is in specific mixes of air pollutants, and how variable the mixes can be from one setting to another. Such information does not seem to be available for any subset of the HAPs.
- k) Allocating the "blame" for bad outcomes that are in fact the result of synergies among exposures (e.g., when we cost out the extra risks and expenses from the synergy between asbestos and tobacco smoke, which industry should get the hypothetical bill?)
- l) The non-linearity of many impairments; for example, one can lose fifty percent of liver capacity and never know it, so that loss of the first ten percent may have a value vastly different from the value of the last ten percent.
- m) Unknown levels of regulatory compliance, and predictably incomplete efforts to

monitor compliance for 189 HAPs

- n) Big and fundamental questions about which things we are to value and about how we are to attach specific values to those things
- o) All the difficulties of trying to place dollar valuations on various kinds of incommensurate health outcomes, including death, and the even greater conceptual difficulties of placing values on markers of exposure or effect when signs or symptoms of illness have not appeared
- p) The sheer volume of 189 HAPs, which will impose great demands for scarce technical talent as well as resources, and which will certainly force the adoption of means to keep those demands within reachable limits
- q) Estimates of the costs of compliance with a new regulation are notoriously prone to error, usually but not always in the direction of gross over-estimation of what it will cost polluters to clean up their act

More fundamentally, different people will value things in different ways. Whose valuation counts? Will we take what people say the first time we ask them, or try to educate them before they give us their values? Will we give special weight to the valuations of people who have had the outcomes in question and understand them? What about substitution effects? What about benefits forgone? Will we assign the same value to every death or every illness of given severity? Willingness to pay is hardly a meaningful metric for someone who has barely enough to get by anyway.

Are dollars even the right metric? There are questions about equity, there are differing and non-linear utilities, there are major questions about what to exclude as externalities. There are discount rates and intergenerational effects to account for, as well as distributional effects more generally, and there may well be important transaction costs.

There are theoretical answers to all or nearly all of these points, but each application of theory requires the use of inputs that are to some extent uncertain. I was quite serious about the six orders of magnitude of uncertainty. We can be honest about that uncertainty, bury our heads in the sand, and see the special interests take over the process, or we can abandon our scientific and technical integrity, lie about the uncertainty, and ultimately lose our credibility and our claim to special standing as scientists. Or, we can come to grips with it as a serious challenge, deal with it directly and honestly, and do what we can to make sure that users of our analyses understand the fundamental falsehood of any claims (including tacit claims) that some other approach is better.

Some recommendations that may be constructive:

- a) There is a need for substantial education about the art of the possible. We need to

educate congress, the public, and the news media; in fact, every group or person who encounters these issues. Perhaps most, we need to educate ourselves, so that we have realistic expectations of what we can in fact accomplish.

b) There is an evident need for very much closer links between risk analysts and cost-benefit analysts. Each person with a significant technical or managerial role on either side of this divide should spend at least six months working in the other program. (When I came to this point in the meeting itself, I heard snickers from the audience. As well as I could tell, they came from a few persons on each side of the divide, though there may be no better way to understand the real problems, and to learn how to help solve those problems, than to wrestle with them yourself. I fear that some persons may not recognize the career advantages of knowing both sides. Heavy pressure from higher levels may be needed to implement this change.)

c) I recommend also that there be regular, weekly meetings on each project in which risk analysis will be a significant element of a cost benefit analysis, to assure full communication and understanding about what is needed, what can be provided, and how to adjust for the inevitable gaps between these. Passing written reports back and forth will not do the job.

d) There is a need for serious attention to the level of accuracy needed at each step of the process of risk analysis / cost benefit analysis. Does it matter if we are off by 20%? Two-fold? Ten-fold? One thousand-fold? Analysts and managers rarely address these questions in any serious way (perhaps because they do not get beyond the correct recognition that greater accuracy is always better, *ceteris paribus*), and yet they are critically important here because of the need to make most effective use of limited resources and to balance countless compromises and trade-offs.

e) We need an organized, almost assembly-line approach to risk assessment if we are to deal with all 189 HAPs. The need for standardized procedures to deal with the HAPs inevitably leads to the need for "bundling", though this too will require some compromises, and general solutions may not always fit well. (Perhaps Procrustes had the right idea.) Bundling according to health endpoint or cause of death might advance the purposes of hazard identification. Bundling according to chemical species could advance exposure estimation. Similarly, dose-response and sensitivity studies might be stronger if we bundle by biologic mechanism or mode of action. Finally, regulatory considerations may fit best with bundling by source categories. The last of these is apparently favored by the economists because it is related directly to their task, but the value of the other axes of classification in the risk assessment phases may outweigh the value of using source categories in the cost-benefit analysis. It may be that we could somehow combine these other axes with source categories to gain some of the advantages of both.

f) The axis of bundling needs more study than has been evident here, since it has implications for the inputs to the cost-benefit analysis, which should follow the same

pattern. If the bundling is by source categories, testing should also be by source categories, with a focus on the study of the complex mix that comes from any one source, rather than its components. I understand the technical and scientific objections to this, but if we take those objections too seriously, they will simply undermine the whole rationale for bundling by source categories. One cannot have it both ways. These matters require serious study before EPA adopts an approach based on source categories as a critical axis of classification.

In summary, we need radical solutions; tinkering with methods, approaches, and mind-sets now on the shelf will not do the job. The technical and scientific issues of cost-benefit analysis of regulatory control of the 189 HAPs are daunting, and any conceivable result will inevitably carry an enormous margin of uncertainty. It is important that everyone involved in this process, including all users of the results, understand that this is true and that it is unavoidable. This is not an indictment of either risk analysis or cost-benefit analysis, both of which are critically important in collecting, analyzing, and interpreting what is or will be on the record. It is certainly better to collect, organize, and interpret what we can than to simply give up and proceed on blind faith that some step is or is not justified. However, we must not expect to produce, or allow others to expect from us, a level of precision and certainty that the process is unable to deliver.

G.3. Economics and Toxicology: Results of a Dialogue on the Prospects for Assessment of Benefits from Regulation of Hazardous Air Pollutants. Comments of Dr. Trudy Cameron, Department of Economics, University of California, Los Angeles, CA.

For some time now, there has been a degree of acrimony between some economists and some toxicologists. This acrimony concerns the nature of the information being supplied by toxicologists to economists for use in the congressionally mandated task of valuing the non-market benefits of environmental regulations. From the economist's point of view, the problem can be characterized as "Why don't they just give us the information we need? Why are they being so uncooperative?" From the toxicologist's point of view, the problem can be characterized as "Why are they asking us for something that is impossible to provide? Why are they being so unreasonable." The recent EPA workshop should reduce this acrimony, and help us focus on the task at hand, by emphasizing the insight that we do not live in a perfect (research) world.

a) What do economists need, in a perfect world, to calculate benefits? People's demand for regulation of hazardous air pollutants is what economists call a "derived demand." People are frequently viewed as demanding environmental regulation not for its own sake, but mainly because the regulation may achieve a reduction in health risks.

Economists are accustomed to dealing with a wide range of derived demands. For example, few people demand electricity for its own sake. Instead, we are willing to pay for electricity because of the services that can be provided by electrical appliances. To understand willingness to pay in a derived demand context, it is helpful to consider the chain of relationships

that form the connection between willingness to pay and the underlying good. To continue the electricity example, consider somebody who uses electricity to heat water. Formal modeling of willingness to pay for a kWh of electricity for water heating will depend on the individual's value of a gallon of hot water and on the efficiency of the water heater (namely, how much electricity is required to produce a gallon of hot water). There are only two functions in this particular chain: (1) how much hot water can be produced from a kWh of electricity, and (2) how much utility (i.e. happiness, satisfaction) the individual gets from a gallon of hot water. Since utility is not directly quantifiable, we measure it by how much money the individual is willing to give up to get that increase in utility.

In the case of derived demand for hazardous air pollutant regulation, there are rather more functions involved in the process of characterizing how a given environmental regulation concerning hazardous air pollutants will ultimately affect individual utilities. (The final step, monetization of regulatory benefits, requires the conversion of a given improvement in individual utility levels into an equivalent income difference. This is the topic of a future workshop, so we will leave this final benefit in terms of utility.)

First, we need to outline some of the constituent functions and identify the argument of each function that is of key interest. Each of these functions will of course be multivariate and subject to uncertainty.

(1)

$Em_j = Em(I_j, \dots)$	Emissions Em from firm j depend upon the firm's inputs and technology, I_j
$A_k = A_k(Em_j, \dots)$	Ambient concentrations in region k depend upon emissions of all contributing firms j
$Ex_i = Ex_i(A_k, \dots)$	Exposure of individual i depends on ambient concentrations in their region k
$Di = Di(Ex_i, \dots)$	Dose received by individual i depends on exposure
$C_i = C_i(Di, \dots)$	Cases of health effects for individual i depends upon the individual's dose
$Si = Si(C_i, \dots)$	Symptoms of individual i depend upon whether they are a victim of health effects
$Ui = Ui(Si, \dots)$	Utility levels are probably most directly influenced by symptoms (compromised function, life expectancy, etc.)

Conceivably, each of these relationships could be studied independently. Since controlled experimental data are rare (and may not reflect the true empirical derivatives in the

field), a considerable amount of modeling will be necessary. By formally modeling all of the factors that determining a particular outcome variable, we take other covariates into account and minimize "omitted variables bias" in the estimated slopes. Specifically:

(2)

$E_{mj} = E_m(I_j, \dots)$ - Emissions of a particular pollutant by a particular firm will depend on the firm's inputs, production technology and abatement efforts, including MACT. It may depend upon factor prices (including the prices of precursors of polluting emissions) and upon the prices of the firm's outputs.

$A_k = A_k(E_{mj}, \dots)$ - Ambient concentrations in a particular region will depend upon the emissions of all firms that contribute to these ambient levels of pollution and upon the "fate and transport" (transfer coefficients) for each firm. Transfer coefficients depend upon weather conditions, season, prevailing winds, the nature of the pollutant, and other factors.

$E_{xi} = E_{xi}(A_k, \dots)$ - Exposure to the pollutant of individual i will depend upon the individual's behavior and patterns of activity, including avoidance behaviors.

$D_i = D_i(E_{xi}, \dots)$ - Dose actually received will depend upon exposure and other factors.

$C_i = C_i(D_i, \dots)$ - Cases as a function of dose level will depend upon the individual's socioeconomic status, current health status, age, gender, and other factors (such as the duration of exposure or cumulative exposure).

$S_i = S_i(C_i, \dots)$ - Symptoms, given that the individual develops a case of the health effect, will depend on the individual's metabolism, access to treatment, age, current health status.

$U_i = U_i(S_i, \dots)$ - Utility may depend only upon the spectrum of symptoms the individual does or does not experience. However, it is possible that utility will be affected even if this particular individual is completely asymptomatic. Perhaps knowledge of exposure creates fear, or the exposure and incidence of cases for other people affects utility. There can be both "use" and "nonuse" demands for relief or prevention of symptoms (including fatal cancers).

Studied in isolation, each of these functions presumably has some explicit approximate mathematical form, the parameters of which must be determined from empirical studies. It is not too far off the mark to suggest that each function in the list above is the province of a distinctly different discipline.

Each of the above functions can be embedded into the next. People do not so much want environmental regulation because they want firms to meet MACT standards. Rather, they want environmental regulation because of what it means for themselves and their families. Thus, in a simple model, utility levels depend directly on symptoms being experienced, and indirectly on each of the contributing factors, all the way back to I_i (which we might interpret as abatement

technology), if the EPA should choose to regulate at that level. Consider one possible characterization of individual utility.

$$(3) \quad U_i = U_i (S_i (C_i (D_i (E_{xi} (A_k (E_{mj} (I_j)))))))))$$

The partial derivatives that we need to know in order to "construct" the effect on utility of a change in abatement technology, I_j appear in the following expression:

			MU_i	MS_i	MC_i	MD_i	ME_{xi}	MA_k	ME_{mj}	
(4)	dU_i	0	&&	&&	&&	&&	&&	&&	&&	dI_j
			MS_i	MC_i	MD_i	ME_{xi}	MA_k	ME_{mj}	M_i	

If the EPA regulates individual firm emissions:

			MU_i	MS_i	MC_i	MD_i	ME_{xi}	MA_k	
(5)	dU_i	0	&&	&&	&&	&&	&&	&&	ME_{mj}
			MS_i	MC_i	MD_i	ME_{xi}	MA_k	ME_{mj}	

If the EPA regulates ambient concentrations:

			MU_i	MS_i	MC_i	MD_i	ME_{xi}	
(6)	dU_i	0	&&	&&	&&	&&	&&	MA_k
			MS_i	MC_i	MD_i	ME_{xi}	MA_k	

In the easiest possible world, each of these partial derivatives would be a nonstochastic scalar. This implies extreme linearity in each respective function. But it is likely that most, if not all, of the constituent functions outlined above are rather nonlinear. It is also possible that the derivative of any one outcome with respect to one cause (say $\partial C_i / \partial D_i$ = the effect of a change in dose of benzene on the incidence of cancer) depends on the doses of other HAPs. There can be interactions between stressors that influence the effect of a change in the level of any one stressor.

The benzene, perchloroethylene, and manganese case studies focus primarily on the MC_i / MD_i link in this very long chain of partial derivatives. But if we are trying to assess the

welfare effects of regulation-induced changes in I_j (or Em_j or A_k), these effects depend upon all of the intervening partial derivatives, not just one of them.

But the utility function proposed above is just one of a number of possibilities. It assumes that individuals derive utility from environmental regulation ONLY insofar as it affects the symptoms they experience. It is an economist's job, however, to try to discern just what it is that contributes to individuals' utility levels. It is possible that an individual's utility level depends on emissions Em_i not only indirectly via the symptoms they experience from the health effects these emissions create, but also directly on emissions levels. Perhaps the utility function looks more like this:

$$(7) \quad U_i = U_i (S_i (C_i (D_i (E_{xi} (A_k (Em_j (I_j))))))) , C_i , D_i , E_{xi} , A_k , Em_j , I_j)$$

If this is the way an individual's utility level is determined, then the individual may derive an increase in utility from a decrease in emissions or ambient concentrations even if this change in ambient concentrations produces absolutely no health effects ($\partial U_i / \partial A_k = 0$)! Individuals are allowed to derive utility from whatever they want. There is no justification for considering only that utility derived directly from health symptoms.

b) Economists prefer to attempt to value the things that enter most directly into people's utility functions. In the environmental regulation context, this usually means symptoms, such as "days of eye irritation," or "days of moderate cough," or even "statistical lives lost." In the chain of partial derivatives outlined above, we would prefer to explore how people are willing to trade off dollars for changes in the level of symptoms. We might ask them directly what they would be willing to pay, or we might ask them to choose among policies that involve different costs and different levels of symptoms. The dollar metric is merely an intermediate device to capture how much of other things they would be willing to give up in order to achieve a reduction in some set of symptoms.

Focusing on the value of symptom changes reduces the dimensionality of the problem in many cases. This is analogous to the way market researchers sometimes reduce the vast number of different automobiles on the market to a much smaller number of attributes. Each auto can be characterized not by its make and model and year, but by the bundle of attributes that it represents (curb weight, acceleration, MPG, age, number of seats, etc.). The advantage of this method is that if we study the market prices of autos as a function of the bundle of attributes each represents, we can infer how market price depends on attributes. Then, if confronted by a new make and model, with a specified set of attributes (preferably within the range of attributes observed in the estimating sample) we can figure out approximately what people would be willing to pay for the new vehicle.

To reduce the 188 HAPs to a smaller set of spanning "symptoms," these symptoms need to be defined rather grossly, of course. Rather than trying to come up with distinct benefits estimates for changes in the level of each of the 188 distinct HAPs, we would instead endeavor to infer the incremental value of changes in each of a smaller set of symptoms. A one-unit

change in the concentration of a particular HAP would then need to be quantified in terms of what that means for what matters to people: the suite of symptoms they are experiencing (e.g. eye irritation, sore throat, cough, fatigue, for some compounds, and more serious endpoints for other compounds). Policy changes with respect to HAPs that result in a bundle of symptom changes will make it necessary to ascertain whether the effects of distinct symptoms on utility are additive, subadditive, or superadditive.

We would like to stock the research "shelf" with a set of estimates concerning people's willingness to pay to avoid increments in each of the set of symptoms that are the usual suspects in HAP assessments. We would also like to be able to say something about interactions among sets of symptoms. If we can build up this inventory, then we have some hope of reconstructing willingness to pay for a reduction in the amount of a particular HAP or set of HAPs that may not have been studied explicitly. This process is known as "benefits transfer." We do not imagine that it would be possible or sensible to conduct a separate economic analysis of willingness to pay to reduce each one of the 188 HAPs. Instead, we would like to study only a few of them explicitly (and bring in results on willingness to pay for symptom reduction from benefits assessments for criterion pollutants or other applications).

Economists have become more and more confident over the years about how and when they can come up with reasonable point and interval estimates for reductions in symptoms. But we must rely entirely on other disciplines to convert a proposed environmental policy into changes in a set of symptoms for some segment of society. What can we do if other disciplines are unable to specify point and interval estimates for the rest of the partial derivatives? This is only a dead end if we restrict utility to be derived only from indicators of health status.

But keep in mind that even if the most likely magnitude of the dose-response function derivative is "zero" at current ambient concentrations, this does not mean that individuals cannot experience a direct increase in utility simply from knowing, for example, that the ambient concentration of a suspected HAP has been reduced.

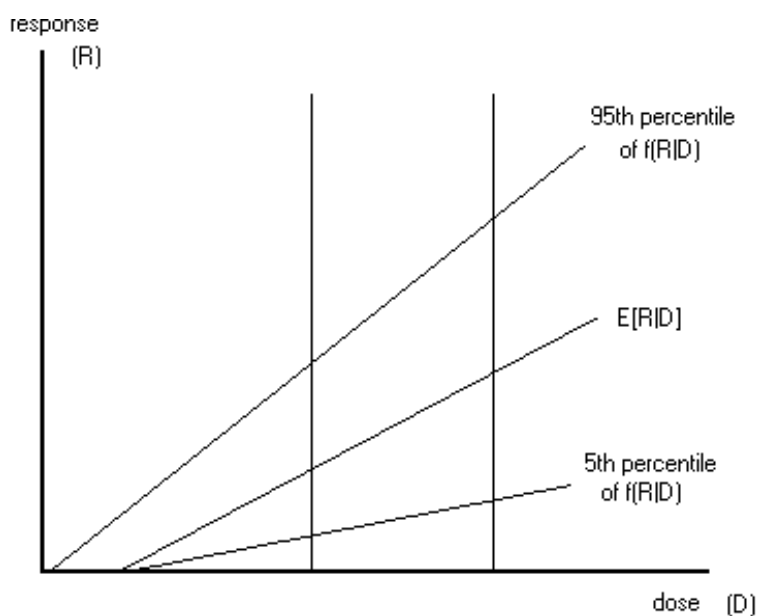
Unlike toxicologists, who rightly expect to be able to identify a mechanism that explains how an increase in the dose of some toxicant contributes to changes in an individual's health status, economists expressly do not seek to figure out how a change in symptoms leads to an effect on individual utility levels. "There is no accounting for tastes." We do not care about WHY people derive utility from something, only that they do. The variant of the utility function that emphasizes only symptoms is consistent with someone having a value system where actions are based on "hypothetical imperatives" in the sense of Kant (i.e. IF you want to achieve this end, then you must take this action). People who care directly about emissions or ambient levels might be viewing HAP control as more of a "categorical imperative" in Kant's terminology (i.e. you must take this action). Individuals can have any of a wide variety of philosophies or value systems driving their individual utility functions. Economists just take these as they are and concentrate on the task of how to aggregate these into a measure of collective social welfare that respects these individual utilities.

c) Why do economists want to know about central tendencies? Can't they figure out what they know from the information about the 95th percentiles???

Since my logistic dose-response curves would be so untidy as to obscure the point, I will draw linear dose-response relationships.

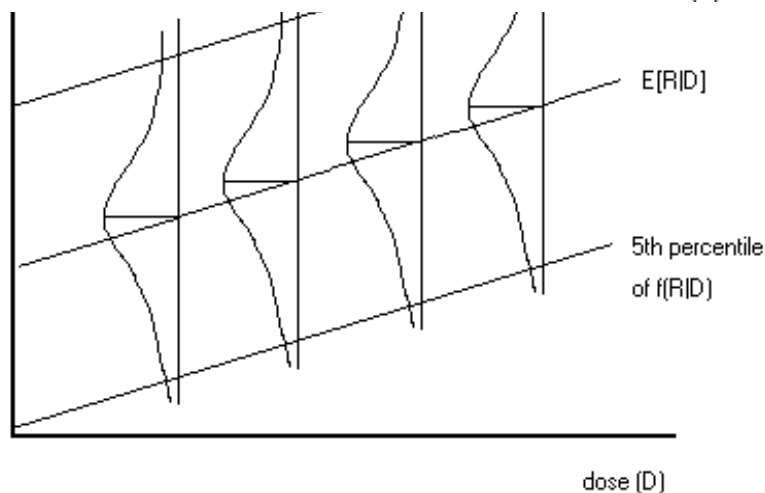
If the conditional distribution of responses at a particular dose had the same shape as the conditional distribution at any other dose (only with a different expected value), then the change in the conditional expectation of the response distribution for a given change in the dose would be identical to the change in the 95th percentile of the response distribution for the same change in the dose. It would not be necessary to know the central tendency.

Are the conditional distributions of responses identical at all dose levels? This is an empirical question. It may be convenient to assume that they are, so that the needed slope (derivative) is



not only constant, but the same at all percentiles of the conditional distribution.

However, without evidence to support this rather heroic assumption, it is more reasonable to allow for a conditional distribution of responses that varies across dose levels. Without drawing the precise shapes of each of these conditional distributions, the lines that connect the 95th percentiles, the expected values, and the 5th percentiles could just as easily look like the second diagram. Here (below), the assumption of linearity in the percentiles is retained, but even this may be untenable.



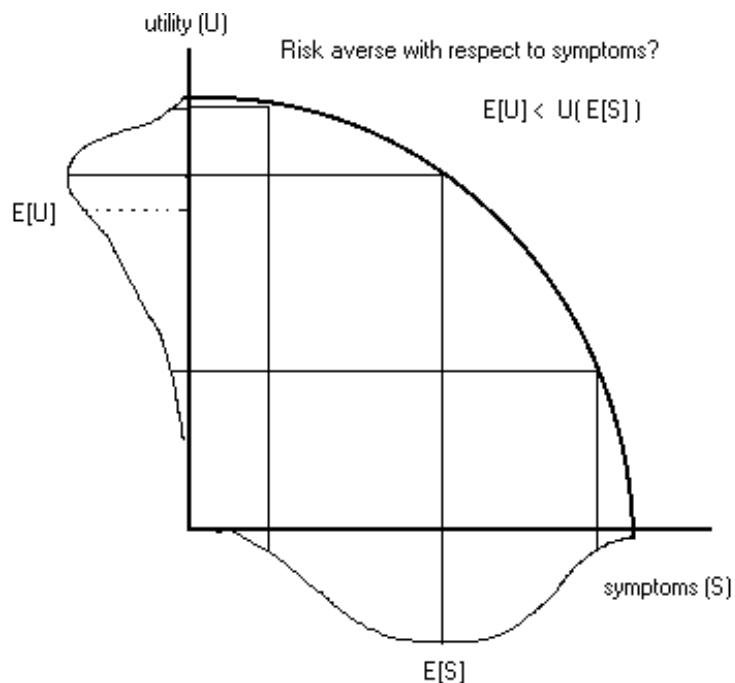
In this more general case, it is clear that the derivative implied by the relationship between expected response and dose will be quite different from the derivative implied by the relationship between the 95th percentile and dose.

Knowledge of the profile of the 95th percentiles of the conditional distribution of response, given dose, is insufficient to determine the complete distribution. Even a normal distribution requires two moments to identify its exact shape. Usually, these are the mean and the variance. But you could equally well pin down a specific normal distribution with information about its 95th and 5th percentiles (because the distribution is symmetric, this implies the location of the mean, and the range between these percentiles implies the dispersion. The 95th percentile alone is insufficient to identify the distribution, even if it was known to be normal (or lognormal.) To be able to identify a distribution based on a single percentile, we need a one-parameter distribution. An example would be the exponential distribution. The location and scale of an exponential distribution would be completely conveyed by the information in its 95th percentile. But can we assert that the conditional distribution of responses, given dose, is exponential. Not without evidence.

d) Why do economists need conditional distributions of response, given dose? If we are going to do a rigorous assessment of the overall effects on benefits from environmental regulation, we need to know both the central tendency and the dispersion in each of the partial derivatives that must be multiplied to produce an estimate of utility gains from HAP regulation. Since the product of random variables is unlikely to have an analytically tractable distribution, simulation methods will typically be required to generate a distribution for benefits that allows the reporting of not only best estimates, but also comprehensive error bars.

e) Utility levels under uncertainty. The problem with a utility function that is linear in an uncertain symptom is that it does not allow for risk aversion with respect to that symptom. A deviation from the expected symptoms that results in lesser symptoms may not have nearly the

negative effect on utility of a deviation in from expected symptoms that results in worse symptoms. The loss of utility when the outcome differs from its expected value may be very much asymmetric on either side of the expected value of the outcome. There has not yet been much empirical work on alternative kinds of utility function for use under uncertainty. Further research on how to implement benefit-cost analysis under uncertainty in more general contexts is clearly needed. Relevant considerations include how individual subjective uncertainty is distinct from scientific uncertainty or disagreement with respect to the ultimate symptom differences to be expected from a HAP regulation.



The associated figure shows a utility function with respect to symptoms, $U(S)$, for an individual who is risk-averse with respect to symptoms. We need to know the distribution of symptoms in order to figure out the distribution of possible utilities associated with these symptoms. Note that a utility function characterized by risk aversion means that a symmetric distribution for S is converted into a skewed distribution for U . The important result is that expected utility $E[U]$ is not simply equal to utility at the expected level of symptoms. It is likely to be less. How much utility is lost due to uncertainty about the level of symptoms depends on the shape of the distribution of S and on the shape of the utility function.

Economists have a pretty good theoretical framework for Benefit-Cost Analysis under uncertainty. A succinct exposition of the theory in the case of objective uncertainty over a binary outcome was presented by Graham (1981). For simple utility functions (e.g. those which are linear in the levels of any factor that is uncertain) it is a straightforward matter to generalize that model to the case of a continuum of possible outcomes (such as the uncertain response to a dose of a HAP, or more directly, the uncertain symptoms from a dose of HAP). When utility is depends upon the level of a symptom AND upon the squared deviation of the symptom from its most likely value, the mathematics is simple and expected utility depends on the expected level of the symptoms and on the variance in symptoms around this expected level.

f) But we do not live in an ideal world. Unfortunately, the perfect ingredients for a full benefit-cost analysis under uncertainty are not available. Some of the key derivatives in the chain are unknown and probably unknowable at finite research cost. How, then, can an economist begin to measure the benefits that people derive from regulation of hazardous air pollutants? Rather than the "bottom-up" approach of building the benefits from derived demand for regulation via demand for reduced health risks, we can consider a more "top-down" approach. This approach accommodates both limited scientific information and individual consumers' subjective assessments of the key partial derivatives (or even just the products of these partial derivatives). I make the following suggestions for a research strategy based on what I am currently attempting to do in the case of derived demand for climate change mitigation policies (Cameron (1998)). Understanding and measuring the demand for expensive climate change mitigation programs has much in common with the problem of understanding and measuring the demand for HAP regulations. Both issues are characterized by uncertainty (incomplete or ambiguous scientific information, controversy among experts, competing corporate and advocacy group positions on the science, varying levels of expert "credibility," and varying degrees of public interest cross-sectionally) and long latency periods.

The individual's subjective distributions on the magnitudes of the partial derivatives in the chain are a product of the interaction between their experiential knowledge and the information that they receive from outside sources. For example, if the individual knows somebody who died of cancer who worked in a dry-cleaning establishment, this knowledge will be combined with whatever expert information to which the individual has been exposed to yield that individual's *subjective* conditional distributions of responses in the dose-response relationship between perc and cancer.

Thus, the proper characterization of individuals' willingness to pay for HAP regulation depends on their subjective assessments of all of the relevant partial derivatives (either individually or in compounded form). An appropriate research strategy would elicit from each individual their *subjective assessments* of what HAP regulation would be likely to achieve in terms of health effects (and other effects). If he or she desires it, the individual should have access to summaries of whatever expert information is available, including the fact that this information is complete, if that is the case. The individual's updated *subjective assessment* should then be established, and then his or her preferred choices among a set of policy alternatives should be elicited. These policy alternatives should differ not only in terms of their costs to the individual, but also in terms of the degree of protection provided against the consequences the respondent expects without regulation. This method accommodates consumer utility from HAP regulation no matter how it arises, whether through avoidance of measurable health effects or simply through "existence" demand for a non-toxic environment. Economists do not presume to question where utility enhancements come from, only whether they exist.

This method of assessing the benefits of HAP regulation is likely to rely upon stated preference techniques (i.e. contingent valuation or its generalizations), since it is hard to imagine actual referenda being held on alternative HAP policies. Fortunately, researchers are

understanding more and more about the limitations and idiosyncracies of stated preference research and how to minimize these.

It is important to keep in mind that economic research concerning the apparent value of HAP regulations conditional on the public's understanding of the risks these compounds present *does not mean that we have to make policy based on widely held misperceptions about the true risks of HAP concentrations*. The idea is to model benefits as an explicit function of *perceived* risks. Once this function is understood, it is then straightforward to replace the *subjective* risks with *scientifically supportable* levels of objective risk and re-calculate the implied level of benefits that would accrue to each individual if their beliefs were consistent with the science. Note that the science can still involve uncertainty, and the values that individuals are "predicted" to hold for HAP regulation should definitely be conditional on the extent of uncertainty (either individual subjective uncertainty, or scientific uncertainty).

g) A smattering of philosophy. Benefit-Cost Analysis, as it is interpreted in most contexts today, is understood to be based upon a utilitarian (Benthamite) social welfare function. It is important to keep in mind that this particular social welfare function is not the only game in town, although most economists in the U.S. are steeped in the utilitarian tradition because it makes the benefit-cost problem tractable and it does have a number of desirable properties. (See Kolstad (2000).)

Imposing an environmental regulation would be a "no-brainer" if it made nobody worse off and at least somebody better off. Then nobody would be opposed to it and at least one person would be in favor. Environmental regulations are controversial only if the beneficiaries gain at someone else's expense. Some individual utilities will go up (or the regulation would not be demanded) but other people's utilities will go down. This is true, for example, if the beneficiaries of the regulation are not the same people that bear the costs.

Suppose there are some winners from regulation and some losers. If the winners win big enough to be able to compensate the losers for their losses, then it would be possible to achieve unanimity about the desirability of a regulation.

Just before WWII, Nicholas Kaldor and John Hicks proposed that the secondary consideration of the distributional consequences of some proposed reallocation of resources can be separated from the primary discussion of whether the net change in utility is positive overall. Even if compensation does NOT take place, if the "benefits" in terms of gained utility for some members of society exceed the "costs" in terms of lost utility for other members of society, then there is an argument that the proposed reallocation is a good idea for society as a whole. It is a second-order issue to then consider whether the distributional consequences of the reallocation are sufficiently undesirable to preclude the reallocation on distributional grounds. This so-called "compensation principle" says that if a proposed reallocation would create more gains than losses, then it is socially desirable, even if no compensation occurs.

A Social Welfare Function (SWF) takes the utility levels of individuals and combines them in some fashion to yield a single-dimensioned scalar number for something called "social welfare." Different value systems lead to different candidates for the SWF.

Utilitarian (Benthamite): $W(u_1, \dots, u_N) = \sum_i S_i u_i, \quad S_i \geq 0$

Aggregate social welfare is a weighted sum of the utility of each individual in society. The weights are positive, but need not be equal. We are usually most interested in welfare *changes* from *resource reallocations*: $\Delta W / \Delta x$ (where x is a determinant of individual utilities, u_i). This is a way of denoting the "net benefits" from resource reallocation " Δ ." If these net benefits are positive, the utilitarian (Kaldor-Hicks compensation principle) opinion would be that the reallocation is a good idea. This is Benefit-Cost analysis, as conventionally practiced.

Egalitarian: $W(u_1, \dots, u_N) = \sum_i u_i - \delta \sum_i [u_i - \min_i(u_i)] \quad \delta > 0.$

In this SWF, society cares about the total amount of utility, $\sum_i u_i$, but also about the degree of inequality. If everybody enjoyed identical utility, the term $\sum_i [u_i - \min_i(u_i)]$ would be zero. The negative sign indicates that social welfare is decreased by departures between individual utilities and the lowest individual utility level. The δ parameter dictates the weight on distributional issues (i.e. inequality).

Rawlsian: $W(u_1, \dots, u_N) = \min_i(u_i)$

A society is only as well off as its least fortunate member.

i) A reminder: Arrow's Impossibility Theorem (AIT). This was one of the discoveries for which Kenneth Arrow won the Nobel Memorial Prize in Economics in 1972 (shared with Sir John Hicks of Kaldor-Hicks fame). The result was part of his Ph.D. dissertation. To paraphrase its useful result: *There is no "ideal" way to combine individual preferences into a social choice mechanism.* Since then, people have spent a lot of time tweaking the conditions, trying to figure out under just what modified conditions you CAN produce a nice tidy theory of social decision-making.

But this impasse cannot stop the frequent need to make policy decisions regarding the allocation of resources in some "best" fashion. In practice, the Kaldor-Hicks compensation principle is often used, with ex post consideration of the severity of the distributional consequences. The AIT just reminds us that this decision is not necessarily the only, or necessarily the best, way of deciding about resource reallocations.

Some Criticisms of the Utilitarian Approach to Decisionmaking:

1) Do we each have a utility function that adequately and consistently represents our preferences, especially over time. Tastes change. Preferences can be manipulated by information campaigns (e.g. advertising, public service announcements, "education").

2). Decisions according to the benefits and costs experienced by current members of society are suspect if not all of the affected individuals are taken into consideration. Specifically, future generations are sometimes not represented. Their tastes may differ from current generations.

Should public policy be based on individual preferences (consumer sovereignty) at all, or on what is morally right? Utilitarianism is a branch of teleological ethics. An alternative ethical system is deontological ethics (deontology= science of duty; moral obligations). Immanuel Kant (*Metaphysics of Morals* (1785) considered that actions can be judged by their intrinsic "rightness" and not by the extent to which they serve to further one's goals or aspirations. Two types of imperatives direct our proper behavior:

- 1) *hypothetical imperatives* (present the practical necessity of a possible action as a means of achieving something else which one desires (or which one may possibly desire; e.g. IF you want that, you must do this!)
- 2) *categorical imperatives*: present an action as of itself objectively necessary, without regard to any other end. e.g. You must do this!

If you subscribe to a deontological ethics, you would not require a benefit cost analysis to justify environmental regulation, you could justify it solely on the basis of an argument that "humans have no business messing up the environment with HAPs." The major problem in using deontological ethics as a basis for policy is that reasonable people can differ in terms of what they judge to be "intrinsically right." These systems can work pretty well in a homogeneous society, but the more heterogeneous the society, the more difficult it is to agree on what constitutes an intrinsically right course of action with respect to policy. We rely on utilitarianism because it makes policy evaluations easier to effectuate. (See Hackett (1998).)

(Some, and perhaps much, of the animosity towards economic welfare analysis stems from misunderstandings about what it IS. There is a vitally important distinction between using Benefit-Cost Analysis to MAKE environmental decisions, versus using Benefit-Cost Analysis to INFORM environmental decisions.)

j) Summary. As an economist who has struggled for quite some time with a number of different problems in valuation of non-market benefits of environmental goods, I really appreciate having had an opportunity to get confirmation from toxicology experts that there are fundamental (and probably unresolvable) gaps in our knowledge about the measurable health consequences of hazardous air pollutants.

This insight means that an unambiguous, *objectively calculated*, bottom-up measure of the social benefits from HAP regulation is unlikely to be forthcoming. But this certainly does not mean that these benefits are zero. People may be willing to allocate society's resources (including their own) to control of hazardous air pollutants simply because there is a possibility that they could have health effects, even if these have not yet been detected (or are unlikely to be detected

any any feasible cost of scientific research). Deciding upon an alternative valuation strategy is the next step.

It seems likely that stated preference research will be the most fruitful way to proceed. Experts in economics and cognitive psychology will have to address the issue of elicitation of subjective probability distributions for health risks (and other risks) that may emanate from ambient levels of hazardous air pollutants. We will need to study how public perceptions of risk can be modified by new scientific results, or by "education" (propaganda) campaigns. People's values for environmental programs depend upon what *they think* the programs are buying them. Their beliefs may or may not be consistent with current scientific understanding. But this does not preclude a strategy of first elicitation, and then simulation, to ascertain what *would have been* the public's value of a specified program had they fully accepted the best current (and possibly incomplete) scientific knowledge.

References

- Cameron, T. A. "Credibility of Information Sources and the Formation of Individuals' Option Prices for Climate Change Mitigation." . Department of Economics, UCLA, June 1998.
- Graham, D. A. "Cost-Benefit Analysis under Uncertainty." American Economic Review 71, (1981): 715-725.
- Hackett, S. C. Environmental and Natural Resources Economics. Armonk, New York: M.E. Sharpe, 1998.
- Kolstad, C. D. Environmental Economics. New York: Oxford University Press, 2000.

G.4. Comments of Ms. Laurie Chestnut, Stratus Consulting, Boulder CO

a) Comments relevant to workshop agenda questions 1 and 2:

Question 1) Proposed approaches for hazard assessments for selected HAPs that would facilitate benefit assessments for those chemicals.

Question 2) Expert discussants' views on whether it is possible to produce a methodology for developing central tendencies and distributions in hazard assessment for HAPs for use in benefits analyses and how that might best be done\

It seems pretty clear from the information presented at the workshop that available toxicology and epidemiology evidence is insufficient to provide dose-response relationships over relevant ranges of exposures for most of the air toxics for which EPA is required to make regulatory decisions. Data based on human studies tend to be at much higher occupational exposures, requiring heroic assumptions when extrapolating to typical community exposures. Data based on animal studies are suggestive of whether or not, and by what mechanism, a chemical may be harmful, but are very difficult to interpret when it comes to quantitative dose

response for humans. Al McGartland gave references to some published papers that look at drawing dose response from information used to determine reference doses for some chemicals. These should be reviewed by the toxicologists to see if the methods are sound and how widely they could be applied. It appears, however, that the information necessary to do a comprehensive quantitative benefits assessment of air toxics reductions, such as is the goal for the 812 assessment, is not going to be available any time soon.

Some questions arose at the workshop regarding if and how economics would use information on variations in dose response for different population groups. Heterogeneity in dose response is a part of the dose response information needed for a benefits assessment. For each dose response function, we need to know to what population it applies. It does not need to be for the general population. There may be different dose response for different age groups, for example. If certain groups are affected significantly differently than others, this should be included in the presentation of the results of the assessment. Similarly, it is important to know whether, for example, there are 100 people facing a change in risk of cancer of 1 in 100 or a million people facing a change in risk of cancer of 1 in a million. The bottom line in both cases may be one expected cancer case, but the risk situation is quite different.

b) Comments relevant to workshop agenda questions 3 and 4

Question 3) How best to identify limitations and uncertainties in both risk assessment methods and economic models.

Question 4) Suggestions and priorities for a research agenda to address identified gaps in available data and methods needed to conduct HAPs-related benefit analyses

In the regulatory decision making context, there is strong motivation to make reasonable use of the available information, even if all the important questions cannot be answered. In this context, ranges of estimates of risk changes, upper or lower bounds on risks, and other information short of a best estimate of dose response, can be utilized to help inform regulatory decisions. Highly variable, uncertain, and inconsistent information, however, is difficult to use in a quantitative benefits assessment unless there is some way to assess how likely it is that each result is accurate. For example, if some studies have found that a given chemical is a carcinogen, and others obtain negative results for the same chemical, we can say that the expected reduction in cancer cases ranges from zero to the amount suggested by the studies that have found an association. Such inconsistencies in results can result in such a large range in benefits estimates that the assessment is not very useful (e.g., saying that the benefits are somewhere between zero and \$10 billion is not usually very useful information for decision makers). However, if some assessment of the likelihood that each of the available results is accurate can be made, then the assessment can be made more informative. Continuing with the same example, perhaps there may be a basis for determining that studies finding a carcinogenic effect are more likely to be correct than studies that have not found carcinogenic effect. Thus, it may be possible to say that there is a 75% chance that the benefits are \$10 billion and a 25% chance that the benefits are zero. This

does not mean that the only benefits number that should be presented is the \$7.5 billion expected value. The range and the probabilities should be presented.

Probability distributions (or some other form of assessment of the likelihood that various results are correct) are also very useful when combining the many steps involved in a benefits assessment. When many ranges of values are all multiplied together, the high and low values can end up being very far apart. A very wide range in the result is not very useful information, especially when the chances that all the low values or all the high values are correct is small. For example, if we take the highest value at each step and multiply them all together, we get a very high value result that is very unlikely to be correct. The ideal information needed to determine probability distributions on results used as inputs for benefits assessment is seldom available, and some professional judgment is inevitable. Sometimes it is better to say we don't know than to rely on pure guesswork. Where that line is is also a matter of judgment. Sometimes the best we can do is some simple sensitivity analyses on key assumptions in the assessment. For example, if this chemical is not a carcinogen, the answer is X; if it is a carcinogen the studies showing an effect suggest the answer is Y.

G.5. Comments of Dr. A. Myrick Freeman, Department of Economics, Bowdoin College, Brunswick, ME

The question posed for us was “whether it is possible to produce a methodology for developing central tendencies and distributions in hazard assessments for HAPs for use in benefits analyses...” I think that in short the answer is “No,” because of both the variety of endpoints of concern (cancer and a variety of non-cancer endpoints) and the variety of sources of data (human epidemiology, long term animal feeding studies for cancer, and other human and animal data for those non-cancer endpoints that have been studied). Rather, different approaches will be necessary for different endpoints and types of data.

For the 20 or so known human carcinogens, best estimates of dose-response (D-R) functions and uncertainty bounds can be obtained using meta-analyses or applying Monte Carlo methods to the available human data. This leaves us with the questions of extrapolation to low doses and the possible existence of thresholds; but these are well known questions.

For the other possible carcinogens, one possible answer is to try to obtain better human data on D-R relationships; but there are the well known problems of estimating human exposures and the typical low power of epidemiology studies. Lacking additional human data, the only recourse is to the animal data. This would mean obtaining maximum likelihood estimates of cancer slope factors and their confidence intervals from the original animal test data. Again there is the question of high to low dose extrapolation as well as the animal to human concordance and extrapolation questions. But I don't see any other way to proceed. I second Lester Lave's suggestion concerning the comparison of animal and human data for those substances for which both kinds of data exist.

Resorting to the animal data for all of the possible HAPs carcinogens that lack human data is probably impractical. It will be important to perform some kind of screening and prioritizing exercise based on, e.g., indications of large volumes of emissions or where human exposures are thought to exceed RfCs. Screening the noncancer HAPs can proceed in a similar fashion. For example, where present human exposure is less than the RfC, the benefits of further reductions in emissions are likely to be zero.

The preceding advice has been based on the assumption that benefits are being defined in the standard way as willingness to pay (WTP) for reductions in the numbers of cases of various types of disease, eg., cancer, obstructive lung disease, etc. An alternative approach to benefit estimation (which might be pursued in parallel with the standard approach) is to investigate different ways to define the commodity to be valued that reflect the ways in which individuals actually think about reducing the risks of environmental disease. For example, economists have used the averting behavior model to analyze data on bottled water purchases as a way of valuing reductions in the risk of waterborne disease. My conjecture is that many individuals view their purchases of bottled water, not as a means of reducing the risk of specific diseases, but as a means of increasing safety more broadly conceived. If there is anything to this conjecture, then individuals might view a broad or comprehensive policy of controlling emissions of HAPs as producing safety or “peace of mind” rather than as yielding reductions in the risks of specific diseases. And if that is the case, then best estimates of reductions in specific risks are not required for benefit estimation. What would be required, however, is a better understanding of the relationship between controls on emissions of HAPs and individuals’ perceptions of safety or “peace of mind.”

G.6. Comments of Dr. Dennis Paustenbach, Exponent, Menlo Park, CA

Response to workshop agenda Question :2: expert discussants' views on whether it is possible to produce a methodology for developing central tendencies and distributions in hazard assessments for HAPs for use in benefits analyses and how that might best be done.

The answer to this question depends on the level of certainty that one needs to satisfy those who will perform and use benefit/cost analyses. Certainly, we can produce cancer risk estimates for animal and human carcinogens using various models. We have used these models in the past in an attempt to rank (in a relative way) the carcinogenic potency of chemicals. However, like most risk assessors, I don't believe that they can accurately predict the actual human response following low exposure to these substances.

I acknowledge that a couple published analyses have suggested that for the genotoxic chemicals, where we have exposure and epidemiology data, there has been a reasonable level of agreement between the number of cases predicted vs the number observed. However, the doses in these studies were those observed in workplaces of nearly 30-40 years ago and these are not even within 100 fold of the doses to which the public is exposed today. In addition, for the cancer estimates from the basic models to be accurate, they should be attached to the "internal" doses estimated by physiologically-based pharmacokinetic (PB-PK) models. This has been best

exemplified in a paper by Reitz et al (1996) which studied vinyl chloride. In short, for virtually all of the animal carcinogens, we are not in a position to use low-dose models to predict the actual cancer risk in humans. We can, for purposes of benefit/cost analyses, use them as a relative index of hazard which could then be "modified" if one wanted to consider other biological factors (like genotoxicity).

Thus, the answer to the question "is it possible to produce a methodology for developing central tendencies and distributions" is yes for the carcinogens. This could be done and, given the abovementioned caveats, it might serve the purposes of the exercise.

With respect to the non-carcinogens, a different approach would be needed as there is assumed to be no risk associated with doses below certain values. One possible approach is to calculate "margins of safety" (MOS) for the non-carcinogens. In this approach, one would take the EPA Reference Dose (RfD) or the Reference Concentration (RfC) and determine the cost associated with achieving doses below these "safe" concentrations. As was mentioned by Dr. Lave, often the public simply wants "to feel safe"....no matter that scientists may give assurances that current conditions pose no significant risk. Assume, for example, background concentrations of formaldehyde in some cities can reach, under certain conditions, about 50 ppt. An airborne concentration generally thought to pose no risk of even transient eye irritation is about 250 ppt. Perhaps, the public wants then never to have concentrations get above 25 ppt. The economists could then provide information to Congress which indicates that the cost of providing a MOS of 10 for formaldehyde is \$100,000,000. A decisionmaker can then compare this cost to achieve an MOS of 10 for formaldehyde to the cost of achieving an MOS of 10 for another non-carcinogenic chemical, for example manganese, and could then weigh the relative importance of the potential adverse effects. In the case of formaldehyde, the threat is transient eye irritation, while for manganese, it could be premature aging of the nervous system. If the costs were similar to achieve an equivalent MOS, then it is likely that the decisionmaker would choose to regulate manganese to rather than formaldehyde given the major differences in toxicity. This approach doesn't equate to a "life save" but should be a perfectly useful metric for both the economist and regulator.

As shown in the above example, it is necessary that the process of benefit-cost analyses for the HAPs be tackled according to the adverse effect of concern. For example, probably about 33% of the chemicals are listed due to their carcinogenicity, 33% are listed because they are systemic (non-carcinogenic) toxicants, while about 33% are irritants. Each might require a slightly different approach. Nonetheless, each has a dose-response curve (or one could be built) and a distribution around the various points on the dose-response curve could be built (for both the carcinogens and non-carcinogens). Over time, the process would almost certainly be modified as more is learned about its strengths and weaknesses.

G.7. Identifying Limitations and Uncertainties in Risk Assessment and Benefit Measurement Methods. Comments of Dr. V. Kerry Smith, Center for Environmental and Resource Economics Policy, Department of Agricultural and Resource Economics, North Carolina State University, Raleigh, NC

This paper is intended as a summary of some of the issues raised in the SAB/EPA Workshop on Benefit Analysis for Policies that reduce exposures to Hazardous Air Pollution (HAP). My specific focus is on the research needed to address limitations and uncertainties in risk assessment (RA) and benefit measurement (BM) methods. The paper is developed in five short sections after this introduction. The first describes a few features of EPA's conventional practices linking RA and BM. The second explains how the issues posed by HAP are different. In the third I define consequential uncertainty and how this concept may offer an approach to help conceptualizing some of the research needed in this area. Section four discusses some activities in risk assessment and benefit measurement where there appear to be opportunities to co-ordinate research. The last discusses the merits of a parallel research strategy.

a) Conventional Practice. The logic linking RA and BM in the evaluation of criteria air pollutants has been unidirectional. For each air pollutant RA estimates a set of health outcomes such as changes in the probability of premature deaths for the general population, or for specific sub-groups (e.g., elderly, asthmatics, etc.), with a well-defined change in each group's exposure to a specific air pollutant (see Chapter 5 and Appendix D of U.S. Environmental Protection Agency [1997] as examples). Monetary values for the health benefits arising measures for unit changes in each health outcome or risk change. For example, if a reduction in the ambient concentration of particulate matter reduces the expected number of days with respiratory illness, then the monetization of the value of this change generally seeks a measure of the benefits of avoiding a day of respiratory illness. These unit benefit measures may not be associated with the specific source of the health effect. As a result, this approach assumes that a day of respiratory illness is essentially the same regardless of what caused it and therefore would have the same economic consequences for affected individuals.

The unidirectional logic is important. It establishes the equivalent of a chain of functions linking emissions to the ultimate health outcome, and, from the economists' perspective, the change in well-being experienced by each individual (see Trudy Cameron's paper for a more detailed elaboration of this logic).

Any description of the limitations and uncertainties in existing practice usually begins by distinguishing uncertainties that arise because many of the components of the chain of relationships linking emissions of pollutants to changes in well-being are stochastic processes. Reductions in the mortality effects of a pollutant, for example, represent changes in the probability of premature death for specific groups of individuals. Because the outcome is a probability change, the framework used in this case to measure the health effect acknowledges that there are many external and internal influences on an individual's probability of dying in a given year. Exposures to specific air pollutants are only one class of these influences. The specific models used to fill in the logical chain may incorporate the inherent stochastic nature of the process or they may deal with expected values, implicitly assuming the relationship can be treated as exact except for measurement errors. This distinction determines whether the analysis acknowledges what one might term "structural uncertainty".

A second important source of uncertainty will be labeled "estimation uncertainty". This source arises because one must estimate each of the functions comprising the logical chain of linkages from the emission to the change in people's well-being. Measurement error, omitted variables, incorrect functional form, etc. can introduce error and therefore uncertainty in the actual policy description used to implement the chain of linkages from emission to benefit. While in practice we cannot separate these two sources of uncertainty, at a conceptual level it is generally important to distinguish them.

An important lesson that has been derived from two decades of research developing RA and BM methods for air quality regulations is that health outcomes identified in the RA process must be capable of being associated with "things" people can, in other contexts, choose. For example, small changes in mortality risk can be selected by people in a number of types of behavior. The most common source of behavioral information used in benefit measurement is in job choices. Increased risks on the job generally lead to compensating differentials in wage rates (see Viscusi 1993)]. Changes in a person's physical condition, e.g., difficulty in breathing, increase in hypertension, do not generally associate with a choice we can observe people making in another context. In these circumstances a new link in the chain must be introduced. That is, the analyst must associate a change in lung functioning with greater incidence of sick days or an increased change of more serious respiratory illness and then evaluate whether there are actual or stated choices that can be related to these outcomes.

This need to augment the chain of linkages is important because to the extent there is discretion in either the intermediate measure of the health characteristic sought or the observable health outcome that results, risk assessment decisions should be made in ways that facilitate making connection to an economic choice. Otherwise, one simply adds to the uncertainty in the measurement process.

Co-ordination between the design of risk assessment and benefit measurement methods is important for another reason as well. Benefit measurement methods rely on observing (or offering potential opportunities) for people to make choices that reveal how they would tradeoff some outcome that can be related to the pollutant of interest and another commodity that can be expressed in monetary terms. The unidirectional logic of most benefit information for air pollution policy assumes that these choices can be observed in another context and transferred to any evaluation of policies involving air pollutants. Expressed in terms of the chain of linkages discussed by Cameron, and illustrated in simple terms in Figure 1, this process allows monetization of the changes in outcomes at step 5 in the process as approximate measures for the value of their consequences for well being identified as step 6 in this logic.

**Figure 1:Chain of Linkages for Environmental Policy Evaluations
for Air Pollutants**

Characterization of:

Step Number:

Emission Rates for Pollutants by Source 9	1
Spatial Diffusion, Atmospheric Interactions, and Other Influences to Ambient Concentration of Pollutants 9	2
Exposure Patterns for Receptors at Each Location 9	3
Physical Response to Exposure to Pollutants 9	4
Health Outcome Resulting from Physical Response 9	5
Effect on Individual Well-being	6

Unfortunately the unidirectional flow is itself an approximation and the very behavioral choice relied upon to make benefit measures for health outcomes can affect the reliability of the one way flow of causation. That is, revealed preference arguments assume to the extent people recognize the effects from steps 2 to 6, they will react to them and try to adapt.¹ The hedonic model assumes they consider site specific amenities in residential location choices. Models of averting and mitigation behavior suggest that other, less costly, responses will also be made. These can include spending less time outside during high pollution times or purchasing central air conditioning, etc. These choices imply the equivalent of feedback loops between steps 6 and 3. The physical responses to some types of pollution may be more easily recognized by people. As a result, in this case the mitigating responses may be more likely and the feedback important. This dimension also will be important to the RA/BM connections with hazardous air pollutants.

b) Three case studies at the SAB/EPA Workshop illustrated the range of possibilities in information likely to be available about hazardous air pollutants. Benzene offered a case with substantial information (compared to other HAP's) on the relationship between exposure (at high doses) and human health outcomes. While there was little basis for evaluating the extrapolation to low dose levels, there also was no basis for arguing modifications to conventional practices.

¹These averting or mitigating behaviors imply continuous adjustment is feasible. To the extent the set of available choice alternatives is finite, then, as Bockstael and McConnell (1999) suggest, people will select the best alternative in the set. This does not imply the choice will be at a point where the marginal value of the amenity underlying the choice equals the marginal cost of adjusting to obtain it.

By contrast, perchlorethylene illustrated a situation where despite substantial data for humans and animals, the strategies for measuring impacts were incomplete due to data limitations. Thus the evidence available could only be regarded as "weak signals" of potentially more important effects. Nonetheless, over the short term there seemed to be little basis for improving the information available and some perception that judgments connecting exposures to health outcomes would need to rely on encoding experts' judgments rather than more formal empirical tests.

The last case, manganese, had most of its information concentrated in describing aspects of step 5 in the linkage chain given in Figure 1. Completing the linkage required judgments to connect physical responses to more conventional health outcomes and to evaluate how economically meaningful choices could be connected to non-traditional outcomes.

An overall implication of the background presentations by both Farland and Lave was that policymakers are unlikely to have the type or the level of detail in information available for the HAPs to be evaluated. Moreover, decisions about regulating them will be made before the research required to implement the conventional RA/BM logic would be available. Three issues emerged implicitly or explicitly in the resulting discussion: screening rules across HAPs were needed to identify the most likely candidates for regulation; the methods used to conduct RA and BM need to consider the treatment of uncertainties in the component elements of policy analyses evaluating practices based on they become "consequential" for decisions; and future research programs should be structured to include parallel research activities, creating pathways for cross-checking findings. The first of these is discussed in comments prepared by Chestnut and Locke; the second is discussed in the next section; and the last is considered in Cameron's comments and in the last section of this paper.

c) Consequential Uncertainty. Regulatory policy based on risk assessment by definition recognizes that a policy is intended to change the stochastic environment in which lay people must make their decisions.² Thus, a policy evaluation of a regulation in this context describes uncertainties people face with and without a regulation. Estimation and structural uncertainties are both reflected in such descriptions.

The treatment of uncertainty becomes consequential when it alters one or more aspects of the criteria influencing policy decisions in a way that would alter a choice. In short, analysis decisions about how to reflect the sources of uncertainty for HAP don't matter if the policy choices would not be affected. Thus, improvements in Monte Carlo simulation or other analytical details that would improve variance estimates for health effects are not consequential, if policy choices are always based on central tendencies and these conclusions would not change with the refinement.

²I will not attempt to discuss in this short paper subtle distinctions sometimes considered in economics (and psychology) between the terms risk and uncertainty.

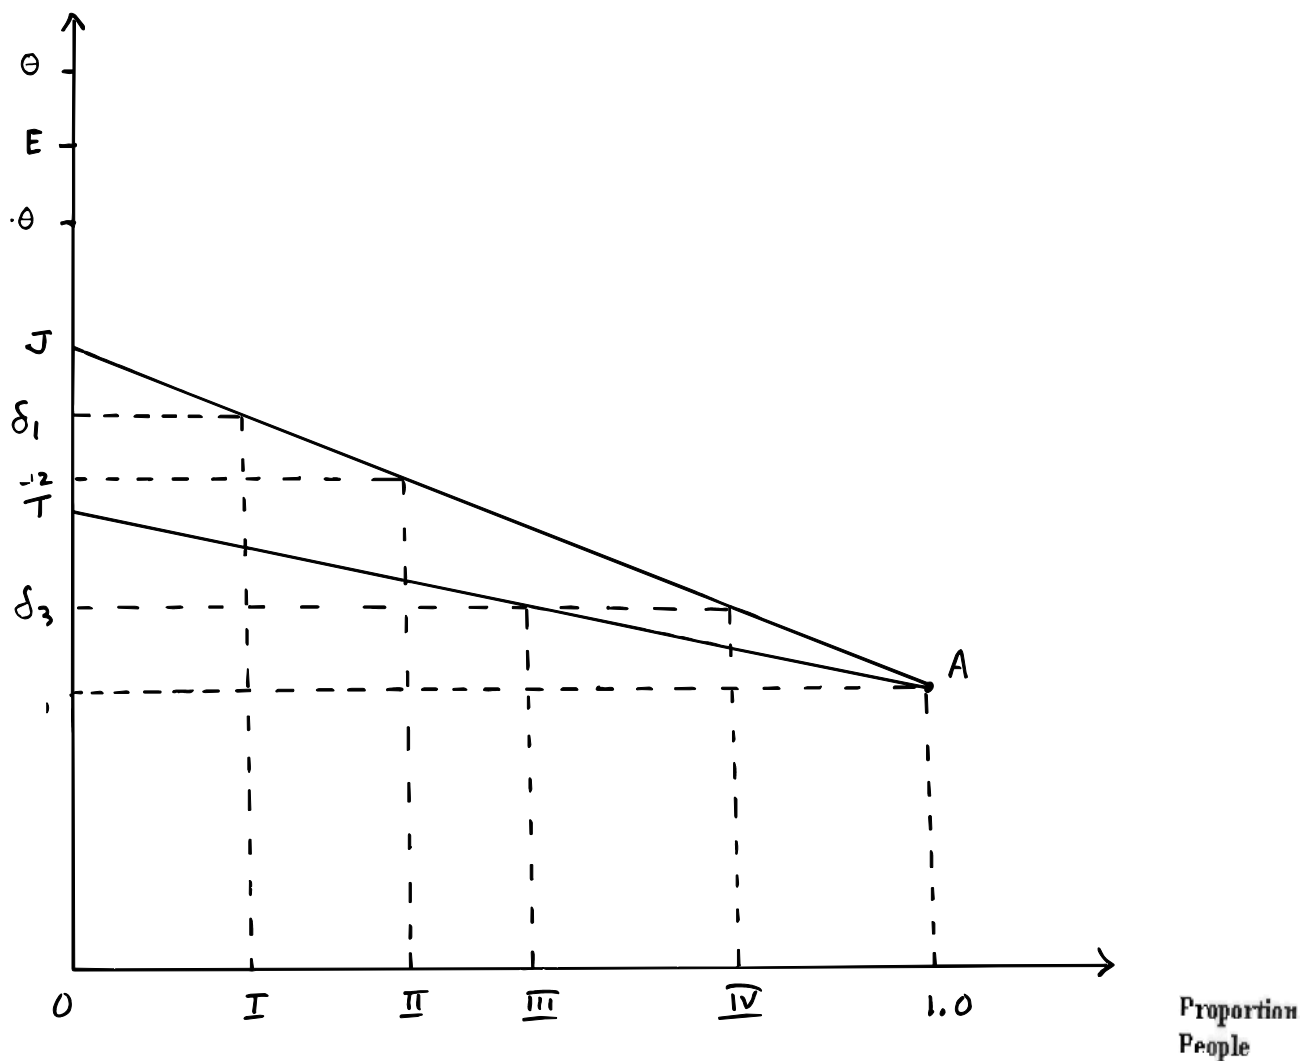
Figure 2 illustrates how the decisions made about coordinating risk assessment and benefit measurement influence whether the treatment of uncertainty is consequential to policy decisions. Five issues contribute to this judgment: the baseline distribution of exposures people receive in the absence of action; the ambient concentration judged to be associated with health outcomes that "count" for regulatory purposes; the estimation uncertainty in describing that ambient concentration; and the risk factor applied based on EPA's propensity to include a margin of safety (due presumably to a composite of concern about structural and estimation uncertainty) complete the elements usually associated with risk assessment. Measures of the economic importance (e.g., unit benefit measures) complete the interacting factors that, together with cost, determine whether regulatory decisions will be consequential to policy choices.

Using Figure 2, on the horizontal axis is plotted the proportion of people experiencing at least the amount of a pollutant measured on the vertical axis.³ The point A0 designates background, so 100% of the population experiences at least this level. For each HAP we would expect a different relationship to describe these exposure profiles. Different population groups could also be represented as having different patterns of exposure. Curves AT and AJ illustrate two alternatives. They are presented as straight lines for graphical convenience only.

³This distribution could refer to the whole population or a specialized group such as children or the elderly.

Feasible
Ambient
Concentration
of Pollutant

Figure 2: Illustration of Effect of Treatment of Components of RA and BM for
“Consequential” Uncertainty



Policy choices to recognize a health outcome (whether traditional or non-traditional) and link it to an ambient concentration implicitly solve "backwards" or invert the functions one associates with the connections implied by the ambient concentration Y exposure level; the exposure level Y physical response; and physical response Y health outcome linkages (steps 2,

3, 4 and 5 in Figure 1). This defines point E and the estimation uncertainty around it in Figure 2. Estimation uncertainty is represented in the figure as the range for $E + S$ to $E - S$. Policy analyses requiring a margin of safety also then impose further risk factors such as a reduction of M in addition to the reduction of S in the concentration regarded as "safe". This specification for a desired (after regulation) ambient concentration determines whether the earlier decisions on how to measure the central tendency and uncertainties in health outcomes will be consequential. That is, consider the cases M , M and M with the two different exposure profiles. My argument says that decisions about whether health outcome is serious enough to count (for regulatory purposes) and about the functions linking ambient concentrations to that outcome implicitly determine the starting point on the vertical axis -- E . Change anyone of them and we move E up or down the vertical axis. Likewise the treatment of estimation uncertainty (q) and judgments about risk factors (S) and judgments about risk factors (i.e., whether M , M or M) will for a given distribution of people (AT versus AJ) yield different proportions of the population that are affected -- 0 I, 0 II, 0 III or 0 IV.

Recognizing these choices as potentially consequential implies we should be evaluating them by asking whether a count of affected people, or a cost per person, or a net aggregate benefit would be important to the regulatory decision. Each is affected in principle by decisions about what counts (and thus the E) and how uncertainty is incorporated (i.e., the S and M). An important element in the discussion at the workshop was that consideration should be given to the implications of making these decisions differently depending on whether the outcomes lead to dramatically different net benefits.

d) Research Complementarities Between Risk Assessment and Benefits Measurement. Table 1 summarizes a few general research tasks that were discussed at the workshop for hazardous air pollutants and the potential for complementarity or interactions that stem from how the research is designed. Three of the four tasks generally defined as involving toxicology or epidemiology would benefit from complementary economic research related to specific modifications in the elements of policy analyses that could be consequential to decisions.

**Table 1 Complementarities in RA⁴ and BM⁵ Research
For Policy Associated with
Hazardous Air Pollutants**

TASK	RA	BM
Evaluating the nature of the physical effect on people	X	-
Evaluating the nature of the health outcome affecting people	X	Eliciting lay person's preferences for different health endpoints
Measuring distribution of ambient concentrations and number of people experiencing them	X	Evaluating the prospects for private action to mitigate or reduce exposure received
Estimating the overall consequences of Baseline (no regulation) and Regulated Alternatives	X	Defining and measuring economic choices for identified health outcomes so tradeoffs could be used to estimate benefits from policy
Evaluating importance of layperson's "worry quotient" or policy as insurance	-	Measure economic benefits as value of a regulatory program or policy

e) Parallel Research. Given the limited information, the number of hazardous air pollutants to be evaluated, and the time and resources available to develop such evaluations Lave's paper and presentation at the workshop suggested a different strategy for measuring the benefits due to HAP regulations.. He proposed that we consider measuring the economic value of the "policy" as an object of choice rather than the reductions in specified health conditions attributed to reduced ambient concentrations of individual hazardous air pollutants.⁶ As noted in Table 1, this approach was discussed as a method for assessing the importance of HAP policy as providing a type of insurance for lay persons' "worries" about serious health outcomes that arise as surprises from exposures to these pollutants.

⁴RA refers to research in field related to risk assessment. The primary areas considered in the workshop discussion were toxicology and epidemiology. An "X" means research is clearly needed.

⁵BM refers to benefit measurement. The elements in the table illustrate research tasks that would be complementary to the task to be addressed in risk assessment

⁶This approach parallels innovations in the use of contingent valuation for the damage assessments associated with natural resource damage cases. To my knowledge it was developed for the Exxon Valdez case by Carson et al. (1992).

The workshop discussions suggested to me an opportunity for economic research to proceed along three inter-related lines that would complement each other and provide opportunities for cross validation of benefit estimates for policy. Unfortunately, there was not sufficient time to discuss this strategy at the workshop so it is not reflected in the summary of this discussion. As a result, this paper ends with a discussion of each of the three lines of activity.

The first entails a proposal implicit in Chestnut and Freeman's comments and in the summary given in Table 1. This encompasses evaluation of whether we can measure people's preferences over non-traditional health outcomes. This task involves not only rating outcomes likely to be associated with hazardous air pollutants, but also investigating the feasibility of using existing revealed preference information and stated preference surveys to recover benefit measures for these types of non-traditional choices. Agee and Crocker's (1994) study of parents' willingness to pay for reducing blood lead levels for their children is an example of the type of revealed preference analysis envisioned in this proposal. Recent applications of conjoint methods (Johnson and Desvousges [1997]) suggest it may be feasible to offer non-traditional health outcomes within this format.

The second line of research involves the focus group, survey development and pilot studies required to evaluate whether Lave's proposals to evaluate the control policy as an object of choice can actually be presented as a plausible choice alternative. It is not clear that it can, but following the protocols used in developing modern CV surveys (especially those for large scale damage assessments) it should be possible to resolve this issue without conducting a full contingent valuation study.⁷

The last line of research was presented briefly in my comments, but time did not permit it to be discussed in specific terms at the workshop. It argues for conducting the first two together because, in principle, we should be able to establish a relationship between measures of the economic value of the policy and the benefits for reducing specific health outcomes. The former is a type of ex ante option price with some private mitigation (as a type of private insurance), and the latter is a set of ex post values for the outcomes being avoided. The early logic developed by Anderson [1979] should, with modification using Graham's [1991] extension to the definition of option price, allow one to relate the two measures under specific conditions. This implies the economic value of the policy (as an object of choice) could be compared to the sum of the economic values of the avoided health outcomes. This would serve to unify the analysis, provide a check for both the CV (i.e. contingent valuation) estimate and a gauge of the potential for omission in cases where only a subset of the physical effects can be measured.

⁷A contingent valuation study rather than conjoint is proposed here because the object of choice is a policy and not a specific set of health outcomes with varying attributes. The full attributes of the results of the policy could not be described. If they could, then more conventional methods would be used. Indeed an identification of the uncertainty in the nature of the avoided health effects would likely be included as part of the description of the policy.

References

Agee, Mark D. and Thomas D. Crocker, 1994, "Parental and Social Valuations of Child Health Information", *Journal of Public Economics*, Vol 55 (September): 89-106.

Anderson, James E., 1979, "On the Measurement of Welfare Cost Under Uncertainty", *Southern Economic Journal*, Vol 45 (December): 1160-1171.

Bockstael, Nancy E. and Kenneth E. McConnell, 1999, "The Behavioral Basis of Non-Market Valuation" in Joseph A. Herriges and Catherine L. Kling, editors, *Valuing Recreation and the Environment* (Cheltenham, U.K.: Edward Elgar).

Carson, Richard T., Robert C. Mitchell, W. Michael Hanemann, Raymond J. Kopp, Stanley Presser and Paul A. Rudd, 1992, "A Contingent Valuation Study of Lost Passive Use Values Resulting from the Exxon Valdez Oil Spill", Report to Attorney General of State of Alaska, NRDA, Inc., November.

Graham, Daniel A., 1992, "Public Expenditures Under Uncertainty: The Net Benefits Criteria", *American Economic Review*, Vol 82 (December): 882-946.

Johnson, F. Reed and William H. Desvousges, 1997, "Estimating Stated Preferences with Rated-Pair Data: Environmental, Health, and Employment Effects of Energy Policies", *Journal of Environmental Economics and Management*, Vol 34 (September): 79-99.

Office of Air and Radiation, 1997, "The Benefits and Costs of the Clean Air Act: 1970 to 1990", U.S. Environmental Protection Agency, October.

Viscusi, W. Kip, 1997, "The Value of Risks to Life and Health", *Journal of Economic Literature*, Vol 31 (December): 1912-1946.